On Differentially Private Filtering for Event Streams

Jerome Le Ny

***Abstract*— Rigorous privacy mechanisms that can cope with dynamic data are required to encourage a wider adoption of large-scale monitoring and decision systems relying on end-user information. A promising approach to develop these mechanisms is to specify quantitative privacy requirements at design time rather than as an afterthought, and to rely on signal processing techniques to achieve satisfying trade-offs between privacy and performance specifications. This paper discusses, from the signal processing point of view, an event stream analysis problem introduced in the database and cryptography literature. A discrete-valued input signal describes the occurrence of events contributed by end-users, and a system is supposed to provide some output signal based on this information, while preserving the privacy of the participants. The notion of privacy adopted here is that of event-level differential privacy, which provides strong privacy guarantees and has important operational advantages. Several mechanisms are described to provide differentially private output signals while minimizing the impact on performance. These mechanisms demonstrate the benefits of leveraging system theoretic techniques to provide privacy guarantees for dynamic systems.**

I. Introduction

Privacy issues associated with emerging large-scale monitoring and decision systems are receiving an increasing amount of attention. Indeed, privacy concerns are already resulting in delays or cancellations in the deployment of smart grids, location-based services, or civilian unmanned aerial systems [1]. In order to encourage the adoption of these systems, which can have important societal benefits, new mechanisms providing clear and rigorous privacy protection guarantees are needed.

Unfortunately, providing such guarantees for a system generally involves sacrificing some level of performance. Evaluating the resulting trade-offs rigorously requires a quantitative definition of privacy, and in the last few years the notion of differential privacy has emerged essentially as a standard specification [2]. Intuitively, a system receiving inputs from end-users is differentially private if one cannot infer from its observable behavior if any specific individual contributed its data or not. Other quantitative notions of privacy have been proposed, e.g., [3], [4], but the differential privacy definition has important operational advantages. In particular, it does not require modeling the available auxiliary information that can be linked to the output of the system of interest to create privacy breaches. Moreover, it is an achievable privacy goal despite the fact that a database on which an individual has no influence could still potentially leak information about her in the presence of arbitrary auxiliary information [2].

Nevertheless, differential privacy is a very strong notion of privacy and might require large perturbations to the published results of an analysis in order to hide the presence of individuals. This is especially true for applications where users continuously contribute data over time, and it is thus important to design advanced mechanisms that can limit the impact on system performance of differential privacy requirements. Previous work on designing differentially private mechanisms for the publication of time-series include [5], [6], but these mechanisms are not causal and hence not suited for real-time applications. The papers [7]–[9] provide real-time mechanisms to approximate a few specific filters transforming user-contributed input event streams into public output streams. For example, [7], [8] consider a private accumulator providing the total number of events that occured in the past. This paper is inspired by this scenario, and builds on our previous work on this problem [10, Section IV] [11, Section VI].

The rest of the paper is organized as follows. Section II provides some technical background on differential privacy and describes a basic mechanism enforcing privacy by injecting white Gaussian noise. Section III describes the real-time event stream filtering scenario of interest. In Section IV, we optimize architectures based on linear estimators to provide real-time private filters with reduced impact on performance. Section V attempts at leveraging the knowledge that the input stream takes values in a discrete set, by considering slightly non-linear structures based on decision-feedback equalization. Finally, we conclude with a brief illustrative example in Section VI.

II. Differential Privacy

In this section we review the notion of differential privacy [12] as well as a basic mechanism that can be used to achieve it when the released data belongs to a finite-dimensional vector space. We refer the reader to the surveys by Dwork, e.g., [2], for additional background on differential privacy, and to [11] for the proofs of the results in this section.

A. Definition

Let us fix some probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Let D be a space of datasets of interest (e.g., a space of data tables, or a signal space). A *mechanism* is a map $M : \mathsf{D} \times \Omega \to \mathsf{R}$, for some measurable output space R, such that for any element $d \in \mathsf{D}$, $M(d, \cdot)$ is a random variable, typically writen simply $M(d)$. A mechanism can be viewed as a probabilistic algorithm to answer a query q, which is a map $q : \mathsf{D} \to \mathsf{R}$.

J. Le Ny is with the department of Electrical Engineering, Ecole Polytechnique de Montreal, QC H3T-1J4, Canada. jerome.le-ny@polymtl.ca

Next, we introduce the definition of differential privacy. Intuitively in the following definition, D is a space of datasets of interest, and we have a symmetric binary relation Adj on D, called adjacency, such that $\mathrm{Adj}(d, d')$ if and only if d and d' differ by the data of a single participant.

Definition 1: Let D be a space equipped with a symmetric binary relation denoted Adj, and let $(\mathsf{R}, \mathcal{M})$ be a measurable space. Let $\epsilon, \delta \geq 0$. A mechanism $M : \mathsf{D} \times \Omega \to \mathsf{R}$ is (ϵ, δ)-differentially private if for all $d, d' \in \mathsf{D}$ such that $\mathrm{Adj}(d, d')$, we have

$$\mathbb{P}(M(d) \in S) \leq e^{\epsilon}\mathbb{P}(M(d') \in S) + \delta, \ \ \forall S \in \mathcal{M}. \quad (1)$$

If $\delta = 0$, the mechanism is said to be ϵ-differentially private.

The definition says that for two adjacent datasets, the distributions over the outputs of the mechanism should be close. The choice of the parameters ϵ, δ is set by the privacy policy. Typically ϵ is taken to be a small constant, e.g., $\epsilon \approx 0.5$ or perhaps even $\ln p$ for some small $p \in \mathbb{N}$. The parameter δ should be kept small as it controls the probability of certain significant losses of privacy, e.g., when a zero probability event for input d' becomes an event with positive probability for input d in (1).

A fundamental property of the notion of differential privacy is that no additional privacy loss can occur by simply manipulating an output that is differentially private. To state it, recall that a probability kernel between two measurable spaces $(\mathsf{R}_1, \mathcal{M}_1)$ and $(\mathsf{R}_2, \mathcal{M}_2)$ is a function $k : \mathsf{R}_1 \times \mathcal{M}_2 \to [0, 1]$ such that $k(\cdot, S)$ is measurable for each $S \in \mathcal{M}_2$ and $k(r, \cdot)$ is a probability measure for each $r \in \mathsf{R}_1$.

Theorem 1 (Resilience to post-processing): Let $M_1 : \mathsf{D} \times \Omega \to (\mathsf{R}_1, \mathcal{M}_1)$ be an (ϵ, δ)-differentially private mechanism. Let $M_2 : \mathsf{D} \times \Omega \to (\mathsf{R}_2, \mathcal{M}_2)$ be another mechanism, such that there exists a probability kernel $k : \mathsf{R}_1 \times \mathcal{M}_2 \to [0, 1]$ verifying

$$\mathbb{P}(M_2(d) \in S | M_1(d)) = k(M_1(d), S), \text{ a.s.}, \quad (2)$$

for all $S \in \mathcal{M}_2$ and $d \in \mathsf{D}$. Then M_2 is (ϵ, δ)-differentially private.

Note that in (2), the kernel k is not allowed to depend on the dataset d. In other words, this condition says that once $M_1(d)$ is known, the distribution of $M_2(d)$ does not further depend on d. The theorem says that a mechanism M_2 accessing a dataset only indirectly via the output of a differentially private mechanism M_1 cannot weaken the privacy guarantee.

B. A Basic Differentially Private Mechanism

A mechanism that throws away all the information in a dataset is obviously private, but not useful, and in general one has to trade off privacy for utility when answering specific queries. We recall below a basic mechanism that can be used to answer queries in a differentially private way. We are only concerned in this section with queries that return numerical answers, i.e., here a query is a map $q : \mathsf{D} \to \mathsf{R}$, where the output space R equals $\mathbb{R}^k$ for some $k > 0$, is equipped with a norm denoted $\|\cdot\|_{\mathsf{R}}$, and the σ-algebra $\mathcal{M}$ on R is taken to be the standard Borel σ-algebra. The following quantity plays an important role in the design of differentially private mechanisms [12].

Definition 2: Let D be a space equipped with an adjacency relation Adj. The sensitivity of a query $q : \mathsf{D} \to \mathsf{R}$ is defined as

$$\Delta_{\mathsf{R}} q := \max_{d,d' : \mathrm{Adj}(d,d')} \|q(d) - q(d')\|_{\mathsf{R}}.$$

In particular, for $\mathsf{R} = \mathbb{R}^k$ equipped with the p-norm $\|x\|_p = \left(\sum_{i=1}^k |x_i|^p\right)^{1/p}$, for $p \in [1, \infty]$, we denote the ℓ_p sensitivity by $\Delta_p q$.

A differentially private mechanism proposed in [13] modifies an answer to a numerical query by adding iid zero-mean Gaussian noise. Recall the definition of the $\mathcal{Q}$-function

$$\mathcal{Q}(x) := \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-\frac{u^2}{2}} du.$$

We have the following theorem [11], [13].

Theorem 2: Let $q : \mathsf{D} \to \mathbb{R}^k$ be a query. Then the Gaussian mechanism $M_q : \mathsf{D} \times \Omega \to \mathbb{R}^k$ defined by $M_q(d) = q(d) + w$, with $w \sim \mathcal{N}\left(0, \sigma^2 I_k\right)$, where $\sigma \geq \frac{\Delta_2 q}{2\epsilon}(K + \sqrt{K^2 + 2\epsilon})$ and $K = \mathcal{Q}^{-1}(\delta)$, is (ϵ, δ)-differentially private.

For the rest of the paper, we define

$$\kappa_{\delta,\epsilon} = \frac{1}{2\epsilon}(K + \sqrt{K^2 + 2\epsilon}),$$

so that the standard deviation σ in Theorem 2 can be written $\sigma(\delta, \epsilon) = \kappa_{\delta,\epsilon}\Delta_2 q$. It can be shown that $\kappa_{\delta,\epsilon}$ behaves roughly as $O(\ln(1/\delta))^{1/2}/\epsilon$. For example, to guarantee (ϵ, δ)-differential privacy with $\epsilon = \ln(2)$ and $\delta = 0.05$, the standard deviation of the Gaussian noise introduced should be about 2.65 times the ℓ_2-sensitivity of q.

III. Filtering Event Streams

We now turn to the description of our scenario of interest, similar to the one introduced in [7], [14]. A system receives an input signal $u = \{u_t\}_{t \geq 0}$ with values in the discrete set $\left\{\pm\frac{k}{2}, k \in \mathbb{N}\right\}$. Such a signal can for example record the number of occurrences of certain events of interest at each period (we centered the values around zero for convenience later on). Similarly to [7], [14], two signals u and u' are adjacent if and only if they differ at a single time by at most d, or equivalently

$$\mathrm{Adj}^d(u, u') \text{ iff } u - u' = k\,\delta_{t_0}, |k| \leq d, \text{for some } t_0, \quad (3)$$

where δ_{t_0} denotes the discrete impulse at t_0. The motivation for this adjacency relation is that a given individual contributes events to the stream at a single time only, and we want to preserve *event-level privacy* [7], that is, hide to some extent the presence or absence of an event at a particular time. This could for example prevent the inference of individual transactions from publicly available collaborative filtering outputs, as in [15].

Even though individual events should be hidden, we would like to release a filtered version Fu of the original signal, where F is a given causal stable linear time-invariant system. Note that in this paper, all signals and filter coefficients are

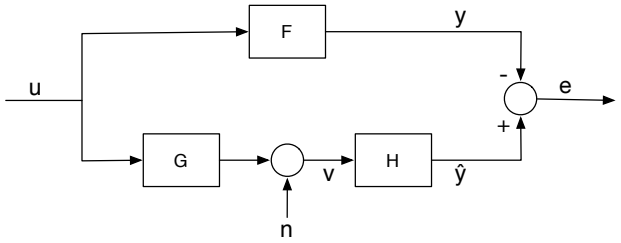

Fig. 1. Differentially private filter approximation set-up. For v to be differentially private, we take n to be a white Gaussian noise with variance $\mathbb{E}[|n_t|^2] = d^2\kappa^2_{\delta,\epsilon}\|G\|_2^2$.

assumed to be real-valued, and all systems are single-input single-output. Privacy preserving approximations of F can be developed based on the following sensitivity calculation.

Lemma 3: Let G be a linear time-invariant system with impulse response $g := \{g_t\}_t$. Then, for the adjacency relation (3) on binary-valued input signals, the ℓ_p sensitivity of G is $\Delta_p G = d\|g\|_p$. In particular for $p = 2$, we have $\Delta_2 G = d\|G\|_2$, where $\|G\|_2$ is the $\mathcal{H}_2$ norm of G.

Proof: For two adjacent binary-valued signals u, u', we have

$$\begin{aligned}\|Gu - Gu'\|_p &= \|G(u-u')\|_p = d\|g * \delta_{t_0}\|_p \\ &= d\|\{g_{t-t_0}\}_t\|_p = d\|g\|_p.\end{aligned}$$

■

This leads to the following theorem, generalizing Theorem 2 to dynamic systems. Certain technical measurability issues in the proof of this result are resolved in [11].

Theorem 4: The mechanism $M(u) = Gu + n$, where n is a Gaussian white noise with covariance $d^2\kappa^2_{\delta,\epsilon}\|G\|_2^2$, is (ϵ, δ)-differentially private for the adjacency relation (3).

Theorem 4 can now be combined with Theorem 1 to obtain a family of privacy preserving mechanisms approximating F, as illustrated on Fig. 1. On that figure, the signal v is differentially private, and hence $\hat{y}$ as well by the resilience to post-processing property (Theorem 1). Two extreme cases include $G = \mathrm{id}$, called input perturbation, and $H = \mathrm{id}$, called output perturbation. In general however, these two choices can exhibit very poor performance [10]. Throughout this paper, we measure the precision of specific approximations by the mean square error (MSE) between the published and desired outputs, i.e.,

$$\lim_{T\to\infty} \frac{1}{T}\sum_{t=0}^{\infty} \mathbb{E}[|e_t|^2],$$

with $e = y - \hat{y}$. The next section is devoted to the description of two ways of choosing the filters G, H as linear filters.

IV. Linear Equalization Mechanisms

A. Linear Zero-Forcing Mechanism

We first recall a mechanism initially described in [10], which we call here the Linear Zero-Forcing (LZF) mechanism. Note that once the differentially private signal $v = Gu + n$ is obtained, the task of estimating y from v is a standard estimation (or equalization) problem. The LZF mechanism is based on the linear zero-forcing equalization idea, and its main advantage is that it requires no statistical information about the input signal u. Let G be a stable, minimum phase filter (hence invertible). Let $H = FG^{-1}$. To guarantee (ϵ, δ)-differential privacy, the noise n is chosen to be white Gaussian with variance $d^2\kappa^2_{\delta,\epsilon}\|G\|_2^2$. The MSE for the LZF mechanism is then

$$\xi^{LZF} = d^2\kappa^2_{\epsilon,\delta}\|G\|_2^2\|FG^{-1}\|_2^2.$$

The best possible choice of filters G is then described in the following theorem [10].

Theorem 5: We have, for any stable, minimum phase system G,

$$\xi^{LZF} \geq d^2\kappa^2_{\epsilon,\delta}\left(\frac{1}{2\pi}\int_{-\pi}^{\pi}|F(e^{j\omega})|d\omega\right)^2.$$

This lower bound on the mean-squared error of the LZF mechanism is attained by letting $|G(e^{j\omega})|^2 = \lambda|F(e^{j\omega})|$ for all $\omega \in [-\pi, \pi)$, where λ is some arbitrary positive number. It can be approached arbitrarily closely by stable, rational, minimum phase transfer functions G.

Note that if $|F(e^{j\omega})|$ satisfies the Paley-Wiener condition

$$\frac{1}{2\pi}\int_{-\pi}^{\pi}\log|F(e^{j\omega})|d\omega > -\infty,$$

then it has a spectral factorization $|F(e^{j\omega})| = \phi^+(\omega)\phi^-(\omega)$ and the bound of Theorem 5 is attained by taking G with impulse reponse

$$g_k = \frac{1}{2\pi}\int_{-\pi}^{\pi}\phi^+(\omega)e^{j\omega k}d\omega, \quad k \geq 0.$$

Note also that the MSE obtained for the best LZF mechanism in Theorem 5 is independent of the input signal u. The design of H does not attempt to minimize the effect of the noise n, as is the case with zero-forcing equalizers [16]. The next section discusses another scheme that achieves a smaller error but requires some additional public knowledge about the statistics of the input signal u.

B. LMMSE Mechanism

The main issue with linear zero-forcing equalizers in communication systems is the noise amplification behavior at frequencies where $|G(e^{j\omega})|$ is small, due to the inversion in $H = FG^{-1}$. However, this issue is not as problematic for the optimal LZF mechanism, since in this case we essentially have $|H(e^{j\omega})| = \sqrt{|F(e^{j\omega})|}$, i.e., the amplification is compensated by the fact that $|F(e^{j\omega})|$ and $|G(e^{j\omega})|$ are both small at the same frequencies. Nonetheless, in this section we explore a scheme based on minimum mean square equalization, which we call the Linear Minimum Mean Square Error (LMMSE) mechanism, and which can exhibit better performance than the LZF mechanism but requires some additional knowledge about the second order statistics of u. This scheme was briefly discussed in [10], but the

optimization of G described below was not performed in that paper.

Hence, assume that that it is publicly known that u is wide-sense stationary with know mean μ and autocorrelation $r_u[k] = \mathbb{E}[u_t u_{t-k}], \forall k$. Without loss of generality, we can then assume μ to be zero, by substracting the known mean of y equal to $F(1)\mu$. The power spectral density of u is denoted P_u, and is assumed to be rational for simplicity.

The LMMSE mechanism is based on designing the filter H as a Wiener filter in order to estimate y from v. For tractability reasons, we derive the performance of the non-causal infinite impulse response Wiener filter, and optimize the choice of G with respect to this choice for H. Once G is fixed, real-time consideration issues can force us to use a suboptimal design with H a causal Wiener filter, or perhaps introducing a small delay.

The non-causal Wiener filter H has the transfer function

$$H(z) = \frac{P_{yv}(z)}{P_v(z)},$$

where P_{yv} is the cross power spectral density of y and v. Since w and u are uncorrelated, we have

$$P_{yv}(z) = P_u(z)F(z)G(z^{-1}).$$

As for P_v, we have, with n a white noise of variance $\sigma^2 = d^2\kappa_{\delta,\epsilon}^2\|G\|_2^2$,

$$P_v(z) = P_u(z)G(z)G(z^{-1}) + \sigma^2.$$

Hence

$$H(z) = \frac{P_u(z)F(z)G(z^{-1})}{P_u(z)G(z)G(z^{-1}) + \kappa_{\delta,\epsilon}^2\|G\|_2^2}. \tag{4}$$

The MSE can then be expressed as

$$\xi^{LMMSE} = \frac{1}{2\pi}\int_{-\pi}^{\pi} \frac{P_u(e^{j\omega})|F(e^{j\omega})|^2}{\frac{P_u(e^{j\omega})}{d^2\kappa_{\delta,\epsilon}^2}\frac{|G(e^{j\omega})|^2}{\|G\|_2^2} + 1} d\omega. \tag{5}$$

Note that we recover the LZF mechanism in the limit $P_u(e^{j\omega}) >> d^2\kappa(\delta,\epsilon)^2$.

1) Privacy-Preserving Filter Optimization: A close-to-optimal filter G for the LMMSE mechanism can then be obtained by optimization, assuming initially that the reconstruction is done with the non-causal Wiener filter H. We discretize (5) at the set of frequencies $\omega_i = \frac{i\pi}{N}, i = 0 \dots N$. Note that all functions in the integral (5) are even functions of ω, hence we can restrict out attention to the interval $[0, \pi]$. We then define the $N+1$ variables

$$x_i = \frac{|G(e^{j\omega_i})|^2}{\|G\|_2^2}, \;\; x_i \geq 0, \tag{6}$$

and the nonnegative constants

$$\alpha_i = P_u(e^{j\omega_i})|F(e^{j\omega_i})|^2, \;\; i = 0, \dots N$$
$$\beta_i = \frac{P_u(e^{j\omega_i})}{d^2\kappa_{\delta,\epsilon}^2}, \;\; i = 0, \dots N.$$

The minimization of the error (5) leads to the following problem (using a trapezoidal approximation of the integrals)

$$\min_{\mathbf{x}} \;\; \frac{1}{2N}\sum_{i=0}^{N-1} \frac{\alpha_i}{\beta_i x_i + 1} + \frac{\alpha_{i+1}}{\beta_{i+1}x_{i+1} + 1} \tag{7}$$

$$\text{s.t.} \;\; \frac{1}{2N}\sum_{i=0}^{N-1} x_i + x_{i+1} = 1 \tag{8}$$
$$x_i \geq 0, \;\; i = 0, \dots N.$$

Note that the constraint (8) comes from the fact that

$$\frac{1}{\pi}\int_0^{\pi} \frac{|G(e^{j\omega})|^2}{\|G\|_2^2} d\omega = \frac{1}{2\pi}\int_{-\pi}^{\pi} \frac{|G(e^{j\omega})|^2}{\|G\|_2^2} d\omega = 1.$$

The optimization problem (7) is convex, and can thus be solved efficiently even for fine discretizations of the interval $[0, \pi]$. The transfer function of the filter G can then be obtained for example by simple interpolation.

Remark 1: Even if the statistical assumptions on u turn out not to be correct, the differential privacy guarantee of the LMMSE mechanism still holds and only its performance is impacted.

2) Causal Mechanism: The previous description of the LMMSE mechanism involves a possibly non-causal filter H. Sometimes, the anti-causal part of this filter might have a fast decreasing impulse response, in which case the scheme can be implemented approximately by introducing a small delay in the release of the output signal $\hat{y}$. Otherwise, we need to implement a causal Wiener filter H. Denoting the spectral factorization of P_v

$$P_v(z) = \gamma_v^2 Q_v(z)Q_v(z^{-1}),$$

we then have

$$H(z) = \frac{1}{\gamma_v^2 Q_v(z)}\left[\frac{P_{yv}(z)}{Q_v(z^{-1})}\right]_+,$$

where, for a linear filter L with impulse reponse $\{l_t\}_{-\infty \leq t \leq \infty}$, $[L(z)]_+$ denotes the causal filter with impulse response $\{l_t \mathbf{1}_{\{t \geq 0\}}\}_t$. Due to the more complex expression for H and the resulting MSE, the design of the optimal filter G in this case is left for future work. Here, we optimize G assuming a possibly non-causal filter H, and then simply modify H afterwards if causality needs to be enforced.

V. Decision-Feedback Mechanisms

In general, solutions to the problem of reconstructing the optimum maximum-likelihood estimator of $\{(Fu)_k\}_{k \geq 0}$ from $\{v_k\}_{k \geq 0}$ are computationally intensive and require the knowledge of the full joint probability distribution of $\{u_k\}_{k \geq 0}$ [16]. This is the main reason why simpler linear architectures such as the one described in Section IV are more often implemented in communication receivers. However, so far, we have not exploited in the estimation procedures the knowledge that the input signal takes discrete values (or perhaps is even binary valued, as in [7], [8]). This can be done by introducing only a slight degree of nonlinearity, using the idea of decision-feedback equalization

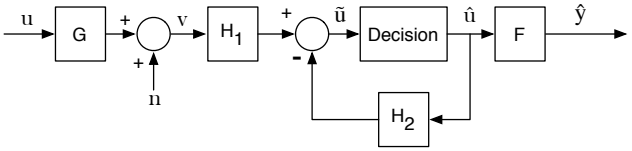

Fig. 2. Decision-feedback mechanism. The decision block is nonlinear and depends on the knowledge about the input signal u, acting as a detector/quantizer.

[16]. We call the resulting mechanism a Decision-Feedback (DF) mechanism. Its architecture is depicted on Fig. 2.

The second stage of a DF mechanism consists of a forward filter H_1, a nonlinear decision procedure (detector or quantizer) to estimate u from $\tilde{u}$, which exploits the fact that u takes discrete values, and a filter H_2 that feeds back the previous symbol decisions to correct the intermediate estimate $\tilde{u}$. H_2 is assumed to be strictly causal, but generally H_1 is taken to be non-causal in standard equalizers, for better performance [17]. Hence, DF mechanisms will typically introduce a small delay in the publication of the output signal $\hat{y}$. In the absence of detailed information about the distribution of u, the decision device can be a simple quantizer for integer valued input sequences, or a detector $\hat{u}_k = \texttt{sign}(\tilde{u}_k)$ for binary valued input sequences.

DF equalizers have a long history, and approximate expressions for their MSE can be derived [17]. For tractability reasons, these derivations invariably make the simplifying assumption that the decisions $\hat{u}$ that enter the feedback filter are correct, i.e., $\hat{u} \equiv u$. Unfortunately, it appears that optimizing G for the resulting approximate expression of the MSE is often not a good strategy, because the simplification results in a filter G that does not need to be adapted to the query F any more (only to P_u). Still, we detail this optimization below and discuss an alternative design strategy for G at the end of the section.

The error between the desired output Fu and the signal $F\tilde{u}$, where $\tilde{u}$ is the input of the detector, is

$$e = F(u - \tilde{u}) = F(u - H_1 v + H_2 \hat{u}),$$

which, under the standard but simplifying assumption that $\hat{u} \equiv u$, gives

$$e \approx F(Bu - H_1 v),$$

with $B(z) = 1 + H_2(z)$ a monic filter (since H_2 is strictly causal). As in section IV-B, minimizing this approximate error (over possibly non-causal filters) requires H_1 to satisfy

$$\begin{aligned} H_1(z) &= B(z)\frac{P_{uv}(z)}{P_v(z)} \\ &= B(z)\frac{P_u(z)G(z^{-1})}{P_u(z)G(z)G(z^{-1}) + d^2\kappa_{\delta,\epsilon}^2\|G\|_2^2}. \end{aligned}$$

For this choice of H_1, the approximate MSE becomes

$$\xi^{DF} \approx \frac{1}{2\pi}\int_{-\pi}^{\pi} \frac{P_u(e^{j\omega})|B(e^{j\omega})|^2|F(e^{j\omega})|^2}{\frac{P_u(e^{j\omega})}{d^2\kappa_{\delta,\epsilon}^2}\frac{|G(e^{j\omega})|^2}{\|G\|_2^2} + 1} d\omega. \quad (9)$$

Assuming now the spectral factorizations

$$\begin{aligned} P_u(e^{j\omega}) &= \gamma_u^2 |Q_u(e^{j\omega})|^2 \\ |F(e^{j\omega})|^2 &= \gamma_F^2 |Q_F(e^{j\omega})|^2 \\ \frac{P_u(e^{j\omega})}{d^2\kappa_{\delta,\epsilon}^2}\frac{|G(e^{j\omega})|^2}{\|G\|_2^2} + 1 &= \gamma^2 |Q(e^{j\omega})|^2, \end{aligned}$$

with Q, Q_u and Q_F canonical filters (monic, causal and minimum-phase), the approximate error (9) can be minimized by setting

$$B(z) = \frac{Q(z)}{Q_u(z)Q_F(z)}.$$

The minimum approximate MSE is then

$$\begin{aligned} \xi^{DF} &\approx \frac{\gamma_u^2\gamma_F^2}{\gamma^2} \quad (10) \\ &\approx \gamma_u^2\gamma_F^2 \exp\left(-\frac{1}{2\pi}\int_{-\pi}^{\pi} \ln\left(\frac{P_u(e^{j\omega})}{d^2\kappa_{\delta,\epsilon}^2}\frac{|G(e^{j\omega})|^2}{\|G\|_2^2} + 1\right) d\omega\right). \end{aligned}$$

The last expression is based on a well-known formula for γ^2, see [18, p.105]. Hence we see that an artifact of this approach is that the influence of F and G is decoupled, and thus the minimization of (10) over G leads to a solution that is independent of F, which is generally undesirable. For example, for u iid with $P_u(e^{j\omega}) \equiv 1$, optimizing (10) gives the trivial solution $G(e^{j\omega}) \equiv 1$, and the whole mechanism reduces to an input perturbation scheme with an additional decision stage. Nonetheless, for completeness we mention that optimizing (10) over the choice of G can be done using a discretization approach similar to the one used in Section IV-B.1, now solving the convex optimization problem

$$\begin{aligned} \max_{\mathbf{x}} \quad & \frac{1}{2N}\sum_{i=0}^{N-1} \ln(\beta_i x_i + 1) + \ln(\beta_{i+1}x_{i+1} + 1) \quad (11) \\ \text{s.t.} \quad & \frac{1}{2N}\sum_{i=0}^{N-1} x_i + x_{i+1} = 1 \\ & x_i \geq 0, \;\; i = 0, \ldots N. \end{aligned}$$

In view of these issues, we mention an alternative design strategy for DF-mechanisms. Note from (4) that the (non-causal) LMMSE mechanism involves a reconstruction filter $H(z) = F(z)H_u(z)$, with H_u the LMMSE estimator for u. Therefore we can interpret the DF mechanism on Fig. 2 as introducing an additional stage to the linear mechanisms, to discretize the estimate of u, and replacing H_u by H_1. A strategy to improve on the performance of the LMMSE (or LZF) mechanism is then to keep the same prefilter G designed in Section IV-B, but simply replace the Wiener filter by a decision-feedback equalizer. Our preliminary results tend to confirm that good performance is achievable with this strategy.

VI. Example

Consider approximating the filter

$$F(z) = \frac{1 + 0.995z^{-1}}{1 - 0.995z^{-1}},$$

with the privacy parameters set to $\epsilon = \ln 3$, $\delta = 0.05$. The (wide-sense stationary) input signal is assumed to be binary valued, i.e., $u_t \in \left\{\pm\frac{1}{2}\right\}$ for all t, with zero mean and power spectral density

$$P_u(z) = \frac{3/4}{\left(1 - \frac{1}{2}z^{-1}\right)\left(1 - \frac{z}{2}\right)}.$$

Such a signal can be generated by a two-state Markov chain in the stationary regime, with transition probability matrix

$$\begin{bmatrix} 3/4 & 1/4 \\ 1/4 & 3/4 \end{bmatrix},$$

one state corresponding to the input $-1/2$, and the other state corresponding to the input $1/2$, see, e.g., [19]. In this context we can imagine that the transitions are generated by individual users, and we want to prevent an adversary analyzing the trace $\{(Fu)_t\}_t$ to infer with confidence in which state the chain was at a particular time.

We designed four mechanism: LZF, LMMSE with G optimized based on (7), DF with G optimized based on (11), and DF with the same G as for the LMMSE mechanism. The DF estimators introduce a 5-period delay in the production of the estimate (finite impulse response equalizers were implemented here, based on [17]). Typical sample paths for these four mechanisms are shown on Fig. 3. The theoretical root MSE (RMSE) for the LZF and (non-causal) LMMSE mechanisms are 8.82 and 7.43 respectively. We see that the DF mechanisms significantly reduces the fluctuations in the produced output. Moreover, the LMMSE pre-filter G leads to a clearly better performance for the DF mechanism than the one based on (11) in this case. The magnitude of the frequency response $|G(e^{j\omega})|$ is shown on Fig. 4 for both filters. The cutoff of the LMMSE pre-filter occurs much earlier, taking into account the fact that F filters the high frequencies of u anyway, and this helps to reduce the degradation due to the privacy-preserving noise n.

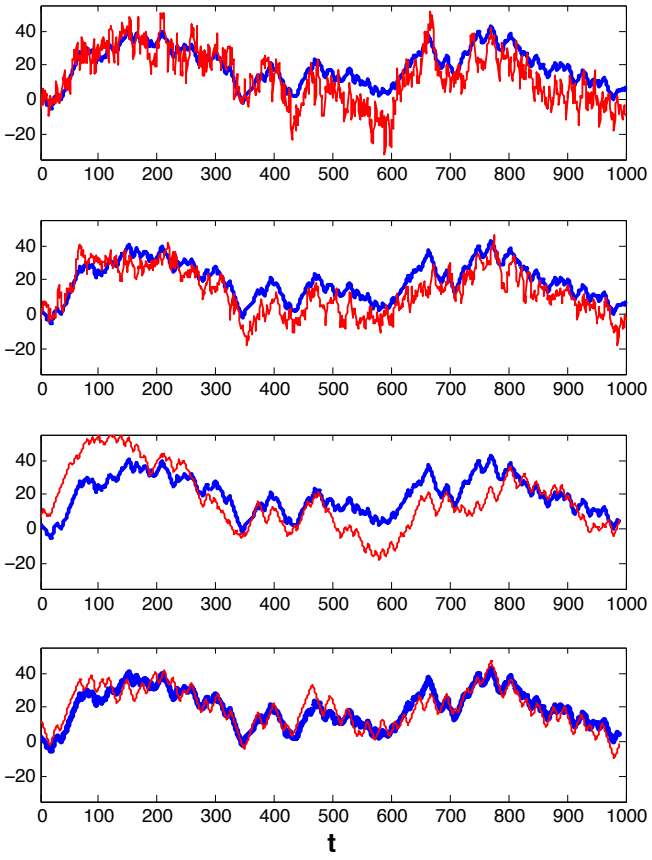

Fig. 3. Sample Paths for four mechanisms. From top to bottom: LZF, LMMSE, DF with G optimized based on (11), and DF with the same G as for the LMMSE mechanism. The original non-private output is shown as the thick blue line.

VII. Conclusions and Future Works

In this paper, we have described several estimation techniques that can be leveraged to minimize the impact on performance of a differential privacy specification for the filtering of event streams. The architecture considered here for the privacy mechanisms decomposes the problem into a standard equalization problem, for which many alternatives techniques could be used, and a first-stage privacy-preserving filter optimization problem. Future work on differentially private filtering for event streams includes enforcing privacy in scenarios where a single end-user can generate events at multiple times, optimizing SIMO and MIMO architectures from a state-space perspective, and adaptive mechanisms that work in the absence of statistics for the input signals.

References

[1] Electronic privacy information center. Online: http://epic.org/.

[2] C. Dwork, "Differential privacy," in *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP)*, ser. Lecture Notes in Computer Science, vol. 4052. Springer-Verlag, 2006.

[3] G. Duncan and D. Lambert, "Disclosure-limited data dissemination," *Journal of the American Statistical Association*, vol. 81, no. 393, pp. 10–28, March 1986.

[4] L. Sankar, S. R. Rajagopalan, and H. V. Poor, "A theory of privacy and utility in databases," Princeton University, Tech. Rep., February 2011.

[5] V. Rastogi and S. Nath, "Differentially private aggregation of distributed time-series with transformation and encryption," in *Proceedings of the ACM Conference on Management of Data (SIGMOD)*, Indianapolis, IN, June 2010.

[6] Y. D. Li, Z. Zhang, M. Winslett, and Y. Yang, "Compressive mechanism: Utilizing sparse representation in differential privacy," in *Proceedings of the 10th annual ACM workshop on Privacy in the electronic society*, October 2011.

[7] C. Dwork, M. Naor, T. Pitassi, and G. N. Rothblum, "Differential privacy under continual observations," in *STOC'10*, Cambridge, MA, June 2010.

[8] T.-H. H. Chan, E. Shi, and D. Song, "Private and continual release of statistics," *ACM Transactions on Information and System Security*, vol. 14, no. 3, pp. 26:1–26:24, November 2011.

[9] J. Bolot, N. Fawaz, S. Muthukrishnan, A. Nikolov, and N. Taft, "Private decayed sum estimation under continual observation," September 2011, http://arxiv.org/abs/1108.6123.

[10] J. Le Ny and G. J. Pappas, "Differentially private filtering," in *Proceedings of the Conference on Decision and Control*, Maui, HI, December 2012.

[11] ——, "Differentially private filtering," September 2012, conditionally accepted for publication in the IEEE Transactions on Automatic Control, available at http://arxiv.org/abs/1207.4305.

[12] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise

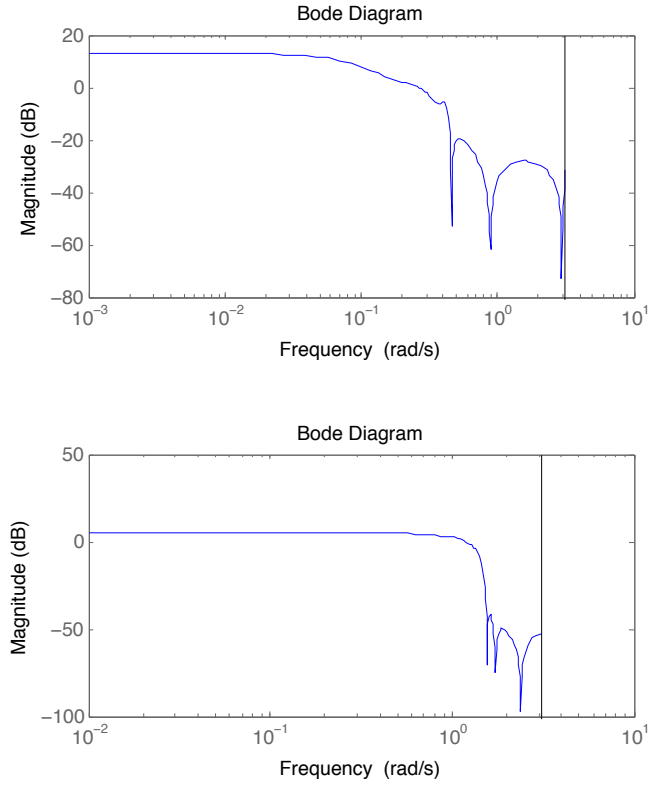

Fig. 4. Magnitude of the frequency reponse $|G(e^{j\omega})|$ for G designed based on (7) (top), and based on (11) (bottom).

to sensitivity in private data analysis," in *Proceedings of the Third Theory of Cryptography Conference*, 2006, pp. 265–284.

[13] C. Dwork, K. Kenthapadi, F. McSherry, I. M. M. Naor, and Naor, "Our data, ourselves: Privacy via distributed noise generation," *Advances in Cryptology-EUROCRYPT 2006*, pp. 486–503, 2006.

[14] T.-H. H. Chan, E. Shi, and D. Song, "Private and continual release of statistics," University of California at Berkeley, Tech. Rep., 2010.

[15] J. A. Calandrino, A. Kilzer, A. Narayanan, E. W. Felten, and V. Shmatikov, ""you might also like": Privacy risks of collaborative filtering," in *IEEE Symposium on Security and Privacy*, Berkeley, CA, May 2011.

[16] J. Proakis, *Digital Communications*, 4th ed. McGraw-Hill, 2000.

[17] P. A. Voois, I. Lee, and J. M. Cioffi, "The effect of decision delay in finite-length decision feedback equalization," *IEEE Transactions on Information Theory*, vol. 42, no. 2, pp. 618–621, March 1996.

[18] M. H. Hayes, *Statistical Digital Signal Processing and Modeling*. Wiley, 1996.

[19] C. Brighenti, B. Wahlberg, and C. Rojas, "Input design using Markov chains for system identification," in *Proceedings of the 48th IEEE Conference on Decision and Control*, Shanghai, China, December 2009.